

11/10/00  
JC950 U.S. PTO

11/13/00

A

PTO/SB/05 (08-00)

Please type a plus sign (+) inside this box → ☐ +

Approved for use through 10/31/2002. OMB 0651-0032

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

JC921 U.S. PTO  
09/10/02

11/10/00

# UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No.	MS154756.1
First Inventor	Garcea, et al.
Title	DISTRIBUTED DATA GATHERING AND ...
Express Mail Label No.	EF185629295US

## APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

- ☒ Fee Transmittal Form (e.g., PTO/SB/17)  
(Submit an original and a duplicate for fee processing)
- ☐ Applicant claims small entity status.  
See 37 CFR 1.27.
- ☒ Specification [Total Pages 29]  
(preferred arrangement set forth below)
  - Descriptive title of the invention
  - Cross Reference to Related Applications
  - Statement Regarding Fed sponsored R & D
  - Reference to sequence listing, a table, or a computer program listing appendix
  - Background of the Invention
  - Brief Summary of the Invention
  - Brief Description of the Drawings (if filed)
  - Detailed Description
  - Claim(s)
  - Abstract of the Disclosure
- ☒ Drawing(s) (35 U.S.C. 113) [Total Sheets 10]
- Oath or Declaration [Total Pages 3]
  - ☒ Newly executed (original or copy)
  - ☐ Copy from a prior application (37 CFR 1.63 (d))  
(for continuation/divisional with Box 17 completed)
    - ☐ **DELETION OF INVENTOR(S)**  
Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
- ☐ Application Data Sheet. See 37 CFR 1.76

ADDRESS TO: Assistant Commissioner for Patents  
Box Patent Application  
Washington, DC 20231

- ☐ CD-ROM or CD-R in duplicate, large table or Computer Program (Appendix)
- Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary)
  - ☐ Computer Readable Form (CRF)
  - Specification Sequence Listing on:
    - ☐ CD-ROM or CD-R (2 copies); or
    - ☐ paper
  - ☐ Statements verifying identity of above copies

## ACCOMPANYING APPLICATION PARTS

- ☒ Assignment Papers (cover sheet & document(s))
- ☐ 37 CFR 3.73(b) Statement (when there is an assignee) ☐ Power of Attorney
- ☐ English Translation Document (if applicable)
- ☐ Information Disclosure Statement (IDS)/PTO-1449 ☐ Copies of IDS Citations
- ☐ Preliminary Amendment
- ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
- ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
- ☒ Other: Express Mail Certificate

17. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment, or in an Application Data Sheet under 37 CFR 1.76:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP)

of prior application No. \_\_\_\_\_/\_\_\_\_\_

Prior application information

Examiner \_\_\_\_\_

Group / Art Unit \_\_\_\_\_

For CONTINUATION OR DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 5b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

## 18. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Label

(Insert Customer No. or Attach bar code label here)

or ☒ Correspondence address below

Name	Himanshu S. Amin				
Address	Amin, Eschweiler & Turocy, LLP				
	24th Floor, National City Center, 1900 East Ninth Street				
City	Cleveland	State	Ohio	Zip Code	44114
Country		Telephone	216-696-8730	Fax	216-696-8731

Name (Print/Type)	Himanshu S. Amin	Registration No. (Attorney/Agent)	40,894
Signature			Date 11/10/00

Burden Hour Statement This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PTO  
09/710172  
11/10/00

# FEE TRANSMITTAL for FY 2001

Patent fees are subject to annual revision

TOTAL AMOUNT OF PAYMENT (\$ 1,306.00

## Complete if Known

Application Number  
Filing Date Herewith  
First Named Inventor Garcea, et al.  
Examiner Name  
Group Art Unit  
Attorney Docket No. MS154756.1

## METHOD OF PAYMENT

1. ☒ The Commissioner is hereby authorized to charge indicated fees and credit any overpayments to:

Deposit Account Number 50-1063  
Deposit Account Name Amin, Eschweiler & Turocy, LLP

- ☒ Charge Any Additional Fee Required Under 37 CFR 1.16 and 1.17  
☐ Applicant claims small entity status See 37 CFR 1.27

2. ☒ Payment Enclosed:

☒ Check ☐ Credit card ☐ Money Order ☐ Other

## FEE CALCULATION

### 1. BASIC FILING FEE

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description	Fee Paid
101 710	201 355	Utility filing fee	710
106 320	206 160	Design filing fee	
107 490	207 245	Plant filing fee	
108 710	208 355	Reissue filing fee	
114 150	214 75	Provisional filing fee	

SUBTOTAL (1) (\$ 710.00

### 2. EXTRA CLAIM FEES

Total Claims 42  
Independent Claims 5  
Multiple Dependent 0

Extra Claims  
-20\*\* = 22  
-3\*\* = 2

Fee from below  
18  
80  
0

Fee Paid  
396  
160  
0

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description
103 18	203 9	Claims in excess of 20
102 80	202 40	Independent claims in excess of 3
104 270	204 135	Multiple dependent claim, if not paid
109 80	209 40	** Reissue independent claims over original patent
110 18	210 9	** Reissue claims in excess of 20 and over original patent

SUBTOTAL (2) (\$ 556.00

\*or number previously paid, if greater, For Reissues, see above

## FEE CALCULATION (continued)

### 3. ADDITIONAL FEES

Large Entity Fee Code (\$)	Small Entity Fee Code (\$)	Fee Description	Fee Paid
105 130	205 65	Surcharge - late filing fee or oath	
127 50	227 25	Surcharge - late provisional filing fee or cover sheet	
139 130	139 130	Non-English specification	
147 2,520	147 2,520	For filing a request for ex parte reexamination	
112 920*	112 920*	Requesting publication of SIR prior to Examiner action	
113 1,840*	113 1,840*	Requesting publication of SIR after Examiner action	
115 110	215 55	Extension for reply within first month	
116 390	216 195	Extension for reply within second month	
117 890	217 445	Extension for reply within third month	
118 1,390	218 695	Extension for reply within fourth month	
128 1,890	228 945	Extension for reply within fifth month	
119 310	219 155	Notice of Appeal	
120 310	220 155	Filing a brief in support of an appeal	
121 270	221 135	Request for oral hearing	
138 1,510	138 1,510	Petition to institute a public use proceeding	
140 110	240 55	Petition to revive - unavoidable	
141 1,240	241 620	Petition to revive - unintentional	
142 1,240	242 620	Utility issue fee (or reissue)	
143 440	243 220	Design issue fee	
144 600	244 300	Plant issue fee	
122 130	122 130	Petitions to the Commissioner	
123 50	123 50	Petitions related to provisional applications	
126 240	126 240	Submission of Information Disclosure Stmt	
581 40	581 40	Recording each patent assignment per property (times number of properties)	40.00
146 710	246 355	Filing a submission after final rejection (37 CFR § 1.129(a))	
149 710	249 355	For each additional invention to be examined (37 CFR § 1.129(b))	
179 710	279 355	Request for Continued Examination (RCE)	
169 900	169 900	Request for expedited examination of a design application	

Other fee (specify) \_\_\_\_\_

\* Reduced by Basic Filing Fee Paid

SUBTOTAL (3) (\$ 40.00

## SUBMITTED BY

Name (Print/Type)	Himanshu S. Amin	Registration No. (Attorney/Agent)	40,984	Telephone	216-696-8730
Signature		Date	11-10-00		

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

Atty. Docket No. MS154756.1

## DISTRIBUTED DATA GATHERING AND AGGREGATION AGENT

by

Federico Garcea, Michael S. Murstein,  
Roger W. Sprague, Alexander M. Sutton,  
Michael W. Thomas and Giedrius Zizys

### CERTIFICATION

I hereby certify that the attached patent application (along with any other paper referred to as being attached or enclosed) is being deposited with the United States Postal Service on this date November 10, 2000, in an envelope as "Express Mail Post Office to Addressee" Mailing Label Number EF185629295US addressed to the: Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

Himanshu S. Amin

(Typed or Printed Name of Person Mailing Paper)



(Signature of Person Mailing Paper)

**Title: DISTRIBUTED DATA GATHERING AND AGGREGATION AGENT**

5

**Technical Field**

The present invention relates generally to computer systems, and more particularly to a system and method for gathering and aggregating operational metrics of a plurality of computers cooperating as an entity wherein the entity may be interfaced collectively as a whole and/or individually. Additionally, the system and method may be employed to gather and aggregate operational metrics of a plurality of entities cooperating as a higher entity where a parent entity may be interfaced directly or as part of an even higher collection of parent entities. The gathering of operational metrics is hierarchical with no predefined limits.

10

**Background of the Invention**

15

20

With the advent of Internet applications, computing system requirements and demands have increased dramatically. Many businesses, for example, have made important investments relating to Internet technology to support growing electronic businesses such as E-Commerce. Since companies are relying on an ever increasing amount of network commerce to support their businesses, computing systems generally have become more complex in order to substantially ensure that servers providing network services never fail. Consequently, system reliability is an important aspect to the modern business model.

25

A first approach for providing powerful and reliable services may be associated with a large multiprocessor system (e.g., mainframe) for managing a server, for example. Since more than one processor may be involved within a large system, services may continue even if one of the plurality of processors fail. Unfortunately, these large systems may be extraordinarily expensive and may be available to only the largest of corporations. A second approach for providing services may involve employing a plurality of lesser expensive systems (e.g., off the shelf PC) individually configured as an array to support the desired service. Although these systems may provide a more economical hardware solution, system

management and administration of individual servers is generally more complex and time consuming.

Currently, management of a plurality of servers is a time intensive and problematic endeavor. For example, managing server content (*e.g.*, software, configuration, data files, components, *etc.*) requires administrators to explicitly distribute (*e.g.*, manually and/or through custom script files) new or updated content and/or configurations (*e.g.*, web server configuration, network settings, *etc.*) across the servers. If a server's content becomes corrupted, an administrator often has no automatic means of correcting the problem. Furthermore, configuration, load-balance adjusting/load balance tool selection, and monitoring generally must be achieved *via* separate applications. Thus, management of the entity (*e.g.*, plurality of computers acting collectively) as a whole generally requires individual configuration of loosely coupled servers whereby errors and time expended are increased.

Presently, there is not a straightforward and efficient system and/or process for providing system wide operational metric data of the collection of servers. Additionally, there is no system and/or process for providing system wide operational metric data of a collection of arrays of servers. Some applications may exist that provide operational metrics of an individual server, however, these applications generally do not provide operational metrics across the logical collection of loosely coupled servers. For example, many times it is important to view information from the collection of servers to determine relevant system-wide performance. Thus, getting a quick response view of pertinent operational metrics (*e.g.*, performance, status, health, events) associated with the plurality of servers may be problematic, however, since each server generally must be searched independently. Downloading all operational metric information from each individual server would overwhelm the network and be extremely cumbersome to an administrator to review all of the operational metric information to find problems or determine a state of the array. Furthermore, the complexity would be substantially increased for a collection of arrays.

### **Summary of the Invention**

The present invention relates to a system and method for gathering and aggregating operational metrics (*e.g.*, performance metrics, system events, health, server state) of a plurality of entities acting as a single entity. For example, the entities may include a plurality of members (*e.g.*, computers, servers, clusters) collectively cooperating as a whole. In accordance with the present invention, an interface may be provided wherein a consistent and unified presentation of metric information of a plurality of the entities as a whole may be obtained from any of the members associated with the entity. The system and method provides for operational metrics of members to be gathered and aggregated to provide a single result set for the entity as a whole, such that entity wide performance can be obtained from a single source or requestor.

In one aspect of the invention, the operational metric data is logged to a data store according to operational metric types. The data can be aggregated across time and then stored to the data store. The data can then be accessed by a gathering and aggregation system for aggregating the data into a single result set across members. Each operational metric type can be provided with an aggregation component adapted to transform and aggregate metric data based on the specific operational metric type. For example, if performance metrics information has been requested for the entity as a whole, a performance aggregation component matches up data point values with respect to time for each member and provides a single result set of aggregated data values to the requestor. The data can be aggregated by performing mathematical operations on each time data point for a particular metric type for each entity that provides this performance data. However, if event metrics information is requested for the entity as a whole, an event coalescing component coalesces event data from each member and provides a single result set to the requestor. Each event is assigned a unique event identifier (*e.g.*, GUID), which uniquely identifies the event. The unique event identifier allows for paging functionality, such that reduced manageable blocks or portions of event data can be provided to the requestor. It is to be appreciated that multiple aggregation components can be plugged into the aggregation system for aggregating different types of metric data with respect to the operation of the entity as a whole.

The following description and the annexed drawings set forth in detail certain illustrative aspects of the invention. These aspects are indicative, however, of but a few of the various ways in which the principles of the invention may be employed and the present invention is intended to include all such aspects and their equivalents. Other advantages and novel features of the invention will become apparent from the following detailed description of the invention when considered in conjunction with the drawings.

### **Brief Description of the Drawings**

Fig. 1 illustrates a schematic block diagram illustrating an operation gathering and aggregation system of an entity in accordance with one aspect of the present invention;

Fig. 2 is a schematic block diagram illustrating an operation and failure management system in accordance with an aspect of the present invention;

Fig. 3 is a schematic block diagram illustrating operation gathering and aggregation of an entity in accordance with one aspect of the present invention;

Fig. 4a is a schematic block diagram illustrating aggregation components of the operation gathering and aggregation system in accordance with one aspect of the present invention;

Fig. 4b is a schematic block diagram illustrating gathering and aggregation of metrics in accordance with one aspect of the present invention;

Fig. 5 is a block diagram illustrating aggregation with respect to time of performance data for different time periods and resolutions residing in a data store in accordance with one aspect of the present invention;

Fig. 6a is a schematic block diagram illustrating gathering and aggregation of events in accordance with one aspect of the present invention;

Fig. 6b is a schematic block diagram illustrating gathering and aggregation of health status in accordance with one aspect of the present invention;

Fig. 7a is a flow diagram illustrating operational metric aggregation in accordance with one aspect of the present invention;

Fig. 7b is a flow diagram illustrating performance metrics aggregation in accordance with one aspect of the present invention;

Fig. 7c is a flow diagram illustrating event metrics aggregation in accordance with one aspect of the present invention;

5 Fig. 7d is a flow diagram illustrating health metrics aggregation in accordance with one aspect of the present invention;

Fig. 7e is a flow diagram illustrating health metrics aggregation in accordance with another aspect of the present invention; and

10 Fig. 8 illustrates a block diagram of a system in accordance with an environment of the present invention.

### **Detailed Description of the Invention**

The present invention is now described with reference to the drawings, wherein like reference numerals are used to refer to like elements throughout. The present invention is described with reference to a system and method for gathering and/or aggregating operational metrics from a plurality of members forming an entity. The members of the entity monitor the operational metrics and log this data locally to a data store. An interface can then request operational metrics data from the members *via* a gathering and aggregation system. The gathering and aggregation system requests and receives operational metrics data from the members based on a requested time period, and in some cases, a particular time resolution and an operational metric. The request can be either for a single member or from all members of the entity. The gathering and aggregation system will then aggregate and format operational metric data for a particular operational metric based on the requested time period and resolution. If the request is for an operational metric for the entity as a whole, the gathering and aggregation system will aggregate or coalesce the member data to provide an overall operational metric data set for the entity. Data that is aggregated refers to data that is manipulated such that a reduced data set or result is provided. Data that is coalesced refers to data that is not manipulated to the extent that the data set or result is reduced. The aggregated or coalesced formatted operational metric data set can then be communicated to an interface,

15

20

25



such as a user interface for displaying the data set. Alternatively, the data set can be accessed by a local or remote process, an external user interface, an external consumer or another member or entity not part of entity from which the data set refers. The gathering and aggregation system can include a plurality of pluggable aggregation components dedicated to aggregating or coalescing a particular operational metric based on the data type of the metric.

In accordance with the present invention, an operational gathering and aggregation system is provided that greatly facilitates management and administration of an entity. The operation gathering and aggregation system interface substantially automates system information retrieval by enabling an application to retrieve the operational metric data of the entity from any of a plurality of systems operatively coupled to the entity. A consistent interface is therefore provided wherein the operation metric data of the entity may be retrieved as if the entity were a singular machine – thereby providing a substantial improvement over conventional systems that may require an administrator to individually retrieve metric data from each machine comprising the entity. Thus, the present invention saves time and administration costs associated with conventional systems. Moreover, system troubleshooting is improved since entity members may be considered as a collective whole (*e.g.*, retrieving system wide performance) and/or individual members may be identified and operated upon.

Although the present example will be discussed with reference to a gathering and aggregation system, it is to be appreciated that information and/or data may be aggregated or coalesced based on the operational metric being gathered. Referring initially to Fig. 1, a system 10 illustrates a particular aspect of the present invention related to an operation gathering and aggregation system for gathering and aggregating or coalescing operational metrics (*e.g.*, performance metrics, system events, system health, system status) of a plurality of members cooperating as an entity. A plurality of members (*e.g.*, computers, servers, machines) for example, computer systems 1 through N (N being an integer) 22a through 22d may be operatively coupled to a network 14 thereby forming an entity 12.

Other sources that may not be part of the entity 12, may also be coupled to the network 14 for retrieving gathered and aggregated data from the entity 12 or for gathering and

aggregating raw metric data from the entity 12 by employing its own gathering and aggregation system. For example, an external consumer of data 26 can connect to one of the computer systems 22 through the network 14 to retrieve raw or aggregated metric data or connect to one of the interfaces 16a through 16d to retrieve raw or aggregated metric data. Additionally, a separate user interface 27 can connect to one of the computer systems 22 through the network 14 to retrieve raw or aggregated metric data or connect to one of the interfaces 16a through 16d to retrieve raw or aggregated metric data. Furthermore, a parent entity 28, parallel entities 29 and/or a child entity 30 can connect to any member of the entity or to the member itself for retrieving and passing metric data between entities for gathering and/or aggregating. In order to request and provide specific gathered and aggregated operation information of the entity 12, a plurality of interfaces (*e.g.*, computer monitor) 16a through 16d may provide output, and an input device (*e.g.*, mouse, keyboard) 24a through 24d may provide input requests to the operation gathering and aggregation system 18a through 18d.

As depicted by the system 10, the interface 16 enables an application or process to retrieve, display or monitor the entity 12 from each member 22a-22d and/or from non-members such as any of the components 26-30. The interface 16 provides a consistent interface for an application or process to measure the operational metrics of the entity 12 as if it was a singular machine. Consequently, the user does not have to administer (*e.g.*, gain access to each machine) and configure (*e.g.*, download new content/software) each machine individually. Thus, time is saved and errors are mitigated. It is noted that the interface 16 generally does not have to run on each computer in the system 10. As will be described in more detail below, full entity operation monitoring may be achieved by interfacing to a single member, for example.

The interface 16 may be served with information provided from each member 22a through 22d employing any of the operation gathering and aggregation systems 18a through 18d. This may be achieved by enabling each member to distribute information to the entity 12. Therefore, the interface 16 may provide aggregated information of the entity as a whole through the operation gathering and aggregation system 18 – in contrast to conventional

systems wherein information of a member may be received and displayed only at the individual member employing an operation monitoring system 20a-20d. For example, computer systems 22a-22d processor performance may be displayed as an aggregation of the output of each member of the entity 12. Any of the interfaces 16a through 16d may be provided with a similar consistent result set. It is noted that the members 22a through 22d may also be entities. For example, some members could also be a collection of members represented by an entity. Thus, the entity 12 may include members that are entities in their own right.

Alternatively, the interface 16 is provided with individual operational metrics from any of the operation gathering and aggregation systems 18a through 18d by requesting this information from that particular operation gathering and aggregation system. Furthermore, entity configurations may be modified from any of the interfaces 16 by enabling the user to provide input to the interface and thereby distribute resultant modifications throughout the entity 12. This may be achieved for example, by providing the input to a single member wherein the single member may then distribute the modified configuration throughout the entity 12. It is to be appreciated that other distribution systems may be provided. For example, rather than have entity operation information centrally distributed and aggregated at the single member, individual members 22a-22d may share a master file (*e.g.*, XML) describing the configuration information of each member.

As illustrated in Fig. 2, performance and failure management may be enabled by generating events 48 for the members 46, logging the events, and monitoring the events either from an entity 32 and/or from a member 46. Events are generally data values reflecting member 46 activity and may be logged into data stores 44a-44c for each member. An operation gathering and aggregation system 34 may then query the data stores 44, and aggregate the information by performing statistical analysis (*e.g.*, summing, averaging, RMS, *etc.* on the member data). For example, Windows Management Infrastructure developed by Microsoft provides an infrastructure to discover information about the system 40 and “subscribe” to various event sources (not shown). The event sources may include entity events such as related to replication of files to members, Windows events such as related to

members, monitors (*e.g.*, Microsoft Health Monitor) such as related to resources such as disk and CPU utilization, and related performance counters (*e.g.*, Microsoft PerfMon).

As an example of aggregation, the operation gathering and aggregation system 34 may acquire events from the data stores 44 (*e.g.*, CPU utilization) and perform an average of the member data relating to CPU utilization and thus provide an average entity CPU utilization to an interface 32. Thus, entity administration, monitoring and troubleshooting is improved over conventional systems by providing a single point of access for an application to administer and monitor entity metrics. It is to be appreciated that events 48 may also be characterized as general purpose interrupts that may be triggered at the occurrence of a predetermined condition. Thus, it is understood that a UNIX and/or other operating system may be similarly configured, for example.

Failure management may be facilitated by including a failure management system 36 (*e.g.*, Windows Health Monitor) which provides the ability to monitor event sources such as system resources (disk, CPU), applications services, performance counters, set rules on the sources (*e.g.*, CPU > 90% for 2 minutes), and take actions when the rule thresholds are triggered. For example, if the above example rule “CPU > 90% for 2 minutes” were exceeded, an application may be notified which could then send an e-mail notice and/or a script file may be generated. Rules provide a system to define metrics that determine whether a member/entity is healthy (status = ok), whether problems may occur soon (status = warning), and/or whether there is a problem (status = critical), for example. Although the failure management system 36 is illustrated as residing on the metric gathering and aggregation system 34, a failure management system may reside on each member 46.

Fig. 3 illustrates a block schematic view of the components employed to provide both a singular member view of operation metrics and an aggregate entity view of operation metrics utilizing the gathering and aggregation system of the present invention. Each member 60 can include a system monitor component 66 adapted to monitor member specific operational metrics and log this information to a data store 64 relating to that particular member. Additionally, an aggregator member 50 can include a system monitor component 66 adapted to monitor aggregator specific operational metrics and log this information to a

data store 64 relating to the aggregator member 50. For example, in the case of performance metrics, the system monitor component 66 periodically retrieves performance data values of different metrics from a performance data source (*e.g.*, WMI). The system monitor component 66 then periodically logs the performance data values in the data store 64 related to that particular member. The counter performance data values can be repeatedly logged based on a predefined time period in respective tables, until the configurations settings are changed. It is to be appreciated that component(s) may reside between the performance monitor component 66 and the data store 64 for setting up communication links, accessing data and/or transforming data.

The system monitor component 66 or some component employed by the system monitor component 66 can then dynamically aggregate or collapse the counter performance data values based on the predefined time period to higher time periods (*e.g.*, 10 seconds, 1 minute, 15 minutes, 1 hour, 1 day) as data time points are increased, so that data is provided for larger periods of time and higher time resolutions (*e.g.*, aggregate or collapse across time). Various mathematical methodologies may be employed to perform such aggregation. For example, for an aggregation from ten seconds to one minute, the performance data values would include six points. The data values of these six points could be aggregated to a minute by taking the average, the minimum, the maximum, the last, the weighted average or some other value of the data values of these six points for supplying the one minute data value.

In the case of member events, the system monitor 66 or a component employed by the system monitor 66 can capture these events when they are generated and log these events to the data store 64. Additionally, health or member status based on predefined rules may be logged to the data store 64 by the system monitor 66 or queried directly by the operation gathering and aggregation member 58. It is to be appreciated that any data type relating to the operation metrics of the aggregator 50 and each member 60 may be logged to the corresponding data store 64 and/or queried directly by the aggregator 50.

An interface 56 can provide a request to the operation gathering and aggregation system 58 for operational data for a particular operation metric over a particular time period based on a single member or based on aggregation or coalescing of the operational metric

over the entire entity. The operation gathering and aggregation system 58 requests this information from the data stores 64 through a query component 49. The query component 49 may include error handling. For example, if a member is not available results are returned from the other members and aggregated appropriately, while an error is returned for the unavailable member, which is not utilized to provide the aggregated results. It is to be appreciated that component(s) may reside between the operation aggregation system 58 and the query component 49 for setting up communication links, accessing data and/or transforming data. The operation metric data is provided to the operation gathering and aggregation system 58 for the particular operation metric that is requested. The operation gathering and aggregation system 58 collapses the data of a given operation metric to fit within a particular time period and resolution to be returned to the interface 56 based on the request. If the operation data is to be returned for the entire entity, the operation aggregation system 58 performs one of an aggregation or coalescing of the data based on the particular operation metric.

For example, for performance metrics the performance data values are aggregated at each data time point to provide a single result set of aggregated data time points for a specified time period and resolution. Again various mathematical methodologies may be employed to perform aggregation with respect to performance data values. For example, for aggregation of four members, the performance data values for each time data point would include four points. The data values of these four points could be aggregated by taking the average, the minimum, the maximum, the last, the weighted average or some other value of the data values of these four points for supplying a single aggregated data value for the entity. For event data, each event will be coalesced into a single event result set over a specified time period including reference to the particular member that the event had occurred. A filter component can be employed to limit the event types to be retrieved by the operation and gathering system 58 and returned to the requestor. For server status, each server's status would be determined and a single status would be returned based on a rule set. For example, if nine servers were operational and one server was not, the entity status would be good. However, if any additional servers stopped operating, this would cause the system to be

return an entity status of critical. Another type of operation is the health of the members. In this situation, the data is hierarchical. This type of data would be aggregated over different levels of data to return a composite result set. The health of the entity can then be determined on an entity based rule set employing the composite result set. Alternatively, metrics can be queried for each member and the metrics aggregated or coalesced. An aggregated member based rule set may then be employed to determine the health of the entity.

It is to be appreciated that not all members will return operation metric data or have operation metric data for a particular point in time. In this situation, the gathering and aggregation system disregards the lack of operation metric data and determines an appropriate aggregated operation metric data on valid data that was returned by the members. Additionally, when more data points are returned to the gathering and aggregation system than requested by the interface, the gathering and aggregation system will interpolate down the data points by calculating the width of the time slice represented by each data point (end time-start time/ data points requested), grouping data points from the result sets and then taking an average or sum as appropriate.

Fig. 4a illustrates an example of components that can form the aggregation system 58 in accordance with one aspect of the invention. The gathering and aggregation system 58 of Fig. 4a includes a number of aggregation components adapted to aggregate data according to different operation data types. The gathering and aggregation system 58 includes a performance aggregation component 70, an event aggregation component 72, a health aggregation component 74, an entity status aggregation component 76 and any additional aggregation components 78 for other operation metric types. It is to be appreciated that different methodologies may be performed in aggregating data of different operation metric types, however, the basic methodology of aggregating operation metric data for retrieving an aggregating a result set of information related to the operation of the entity as a whole is the same.

It is to be appreciated that the type of data collected by the performance aggregation component 70, the entity status aggregation component 76, the health aggregation component 74 and any additional aggregation components 78, alternatively can be collapsed into events

at each member and collected by the event aggregation component 72. Additionally, data collected and aggregated by the performance aggregation component 70, the entity status aggregation component 76, the health aggregation component 74 and any additional aggregation components 78 can be collapsed into events for the entity by the event aggregation component 72. For example, if CPU utilization reaches over 90% for the entire entity (*e.g.*, based on collected performance data), an event can be fired providing information of this condition, which can be coalesced into other events by the event aggregation component 72. Furthermore, if memory utilization of N number of systems remains below a certain level for a given period of time (*e.g.*, based on a rule set of the health monitor component 74), an event can be fired providing this information, which can be coalesced into other events by the event aggregation component 72. It is to be appreciated that substantially any metric type can be collapsed into an event at the member level and/or the entity level.

Fig. 4b illustrates a block schematic diagram of the operation of gathering and aggregating operation metric data employing a metric aggregation system 82. A metric monitor component 80 employs a time aggregation or collapsing component to aggregate or collapse metric data based on a specific time period prior to storing the metric data to the data store 64. The metric aggregation component 82 then aggregates or coalesces the metric data for the plurality of members forming the entity by employing an entity aggregation component 83. Therefore, metric data can be aggregated or collapsed over time and then aggregated or coalesced over members to provide a manageable result set for the entity as a whole.

Referring to Fig. 4b and Fig. 5, an example will be discussed with respect to applying the components of Fig. 4b to the gathering and aggregation of performance metric data. The metric monitor component 80 logs performance metric data periodically based on a predefined time interval to the data store 64. The performance metric data is stored in separate predefined time periods for each metric. The performance metric data stored for each metric can be based on a time period defined by a timer event (not shown). The member time aggregation component 81 dynamically collapses or aggregates performance



metric data to larger time periods and larger time resolutions from a first stored period containing a resolution based on the predefined time interval of the timer event. For example, Fig. 5 illustrates a number of stored time periods residing in the data store 64. The metric monitor component 80 logs a metric list 90, a member list 110 and ten second performance data stored for each metric being logged. The ten second performance metric data is stored for metric #1 92A, metric #2 94A, metric #3 96A up to metric #N 100A. The ten second metric performance data includes performance metric data logged every ten seconds defined by the event timer. The time aggregation component 81 then dynamically updates performance metric data for data of larger time periods and resolutions employing the ten second tables. Fig. 5 illustrates that the ten second data is aggregated up to one minute performance metric data 92B, 94B, 96B up to 100B, which is then aggregated to additional performance metric data, all the way up to one day performance metric data 92N, 94N, 96N up to 100N.

Referring again to Fig. 4b, the metric aggregation system 82 will receive a request from an interface or an internal or external source to gather performance information on a metric over a certain time period for either a particular member or for the entity as a whole. The metric aggregation system 82 will then access or query a particular performance metric data time resolution relating to the time period to be displayed for that metric over a single member or over all members. If the request is for performance metric data for the entity, the entity aggregation component 83 will aggregate the metric over the members to find a single performance value for a range of data points over a particular resolution. The aggregated values will then be transformed to appropriate data points for the particular time period and resolution requested. A result set of the aggregated and transformed values will then be transmitted back to the interface or source.

Fig. 6a illustrates a block schematic diagram of the operation of gathering and aggregation as it relates to event monitoring and coalescing of event data employing the metric aggregation system 82. The metric monitor system 80 employs an events monitor component 120 to log event data for each specified event type in an event table in the data store 64. The events to be logged are determined by an event configuration. The event data

can include a timestamp, a unique identification number (GUID) specific to that event, a member at which the event occurred and data specific to that event. Message templates relating to the event are stored in a message table in the data store 64. The event data is mapped to fields of the event tables and message tables utilizing an event mapping component 122. When an event is retrieved by the metric aggregation system 82, data from the events table and the message table are joined and the message template is filled in (*e.g.*, by insertion strings) from properties in the data field from the events table. The filled in message templates are returned to the event aggregation system 82 in the form of an event short message and an event long message. An event short message relates to a short description of the event, while an event long message relates to a longer more detailed description of the event.

The metric aggregation system 82 will receive a request from a requestor (*e.g.*, an interface or a source) to receive event information over a certain time period for either a particular member or for the entity as a whole. The metric aggregation systems 82 can then access or query event information relating to the time period to be received for a single member or over all members. The metric aggregation system 82 includes an event entity coalescing component 124 adapted to coalesce event data into a single event result set for a particular time period requested by the interface. A filter component 126 can be employed to limit the event types to be retrieved by the metric aggregation systems 82. The GUID represents a bookmark into the coalesced unified result set. Therefore, the interface and a paging component 128 can maintain a record of the first and last event of the received coalesced unified result set. The GUID can then be employed to retrieve subsequent and previous portions of the coalesced result set acting as a virtual bookmark into the result set. To retrieve details on a single event the interface can make a second query specifying a specific event GUID which can identify the source for which the event occurred. The event entity coalescing component 124 will retrieve information from the events message table for that specific event GUID and member and return the information to the requestor.

Fig. 6b illustrates a block schematic diagram of the operation of gathering and aggregation as it relates to health monitoring and health data employing the metric

aggregation system 82. The metric monitor system monitors metric data and one or more member health states 133 are determined based on the metric data and a member health rule set 132. The member health states 133 are a compressed hierarchy of states based on a set of rules. The metric aggregation system 82 will then poll one or more health states from the plurality of members. The metric aggregation system 82 includes a health entity aggregation component 134 adapted to aggregate the top level health status of each member or the status of the various health states for a single member. The health entity aggregation component 134 can then determine the health of the entity based on a health entity aggregation rule set 136. Alternatively, the health entity aggregation component 134 can poll health, state or performance metrics directly from the members or the data store 64 of the members, aggregate the metrics and determine an aggregated health of the entity based on a health member aggregation rule set 138.

Fig. 7a illustrates one particular methodology for aggregation of operation data employing the operation gathering and aggregation system 58 of the present invention. In step 200, the operation aggregation system 58 receives a request for operation data from the interface 56. In step 210, the operation gathering and aggregation system 58 builds a query based on parameters received from the interface 56 and passes the query to the members 60. The operation aggregation system 58 then receives the results on the query from the members 60 in step 220 in the form of an array of record sets. The operation gathering aggregation system 58 then aggregates and formats the results for the interface in step 230. For example, the operation aggregation system 58 gathers the record sets from each member into a single result set of the entity, applies necessary transformation (*e.g.*, average across members), interpolation (*e.g.*, average 600 data points to 100) and provides sorting (*e.g.*, order by time). The aggregated and formatted results are then returned to the interface.

Fig. 7b illustrates one particular methodology for gathering and aggregation of performance data employing the operation gathering and aggregation system 58 of the present invention. In step 260, the gathering and aggregation system 58 receives a request from a source (*e.g.*, an internal or external process, an external consumer, a user interface, another entity) for a performance metric over a particular time period. In step 270, the operation

aggregation system 58 queries the members for specific time period data sets for the particular time period requested. In step 280, the operation gathering and aggregation system 58 aggregates the performance data values for each time point across each member within a specified time resolution to obtain a single result set for the entire entity. In step 290, the single result set is returned to the source.

Fig. 7c illustrates one particular methodology for gathering and coalescing of event data employing the operation gathering and aggregation system 58 of the present invention. In step 300, the gathering and aggregation system 58 receives a request from the interface 56 for event metrics over a particular time period. In step 305, the operation aggregation system 58 queries the members for specific time period event data for the particular time period requested. In step 310, the event data from multiple members 60 falling within the time period requested by the interface 56 is coalesced and stored temporarily in a virtual store. In step 315, the operation gathering and aggregation system 58 retrieves and passes a block or portion of coalesced event data to the interface 56. In step 320, the operation gathering and aggregation system 58 monitors whether or not a second request for an additional block of data or details on a specific event have been received from the interface 56. If the operation gathering and aggregation system 58 does not receive a second request (NO), the gathering and aggregation system 58 continues monitoring for a second request in step 320. If the operation gathering and aggregation system 58 does receive a second request (YES), the gathering and aggregation system 58 retrieves and passes an additional block of event data using a first or last event identifier of the previously received block of data or retrieves and passes event specific data using a specific event identifier in step 325. The event identifier provides for data from multiple sources to be coalesced into a virtual data set in addition to a method for acquiring more data relative to an event result set or a particular event in a virtual data set.

Fig. 7d illustrates one particular methodology for gathering and aggregation employing the operation gathering and aggregation system 58 as it relates to health monitoring. In step 350, the gathering and aggregation system 58 receives a request from a source (*e.g.*, an internal or external process, an external consumer, a user interface, another

entity) for health status of the entity. In step 355, the operation aggregation system 58 queries the members for health status. In step 360, the operation gathering and aggregation system 58 applies the entity health aggregation rule set 136 to determine a health state of the entity. In step 380, the operation gathering and aggregation system 58 then returns the entity health status to the source.

Fig. 7e illustrates another particular methodology for gathering and aggregation employing the operation gathering and aggregation system 58 as it relates to health monitoring. In step 370, the gathering and aggregation system 58 receives a request from a source (*e.g.*, an internal or external process, an external consumer, a user interface, another entity) for health status of the entity. In step 375, the operation aggregation system 58 queries the members for member metric data. In step 380, the operation gathering and aggregation system 58 aggregates the metric data of the members. In step 385, the operation gathering and aggregation system 58 applies the member health aggregation rule set 138 to determine a health state of the entity. In step 390, the operation gathering and aggregation system 58 then returns the entity health status to the source.

In order to provide a context for the various aspects of the invention, Fig. 8 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the various aspects of the present invention may be implemented. While the invention has been described above in the general context of computer-executable instructions of a computer program that runs on a computer and/or computers, those skilled in the art will recognize that the invention also may be implemented in combination with other program modules. Generally, program modules include routines, programs, components, data structures, *etc.* that perform particular tasks and/or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the inventive methods may be practiced with other computer system configurations, including single-processor or multiprocessor computer systems, minicomputers, mainframe computers, as well as personal computers, hand-held computing devices, microprocessor-based or programmable consumer electronics, and the like. The illustrated aspects of the invention may also be practiced in distributed computing environments where tasks are performed by

remote processing devices that are linked through a communications network. However, some, if not all aspects of the invention can be practiced on stand-alone computers. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

5           With reference to Fig. 8, an exemplary system for implementing the various aspects of the invention includes a conventional computer 420, including a processing unit 421, a system memory 422, and a system bus 423 that couples various system components including the system memory to the processing unit 421. The processing unit may be any of various commercially available processors, including but not limited to Intel x86, Pentium and  
10       compatible microprocessors from Intel and others, including Cyrix, AMD and Nexgen; Alpha from Digital; MIPS from MIPS Technology, NEC, IDT, Siemens, and others; and the PowerPC from IBM and Motorola. Dual microprocessors and other multi-processor architectures also may be employed as the processing unit 421.

15           The system bus may be any of several types of bus structure including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of conventional bus architectures such as PCI, VESA, Microchannel, ISA and EISA, to name a few. The system memory includes read only memory (ROM) 424 and random access memory (RAM) 425. A basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within the server computer 420, such as during start-up, is  
20       stored in ROM 424.

25           The computer 420 further includes a hard disk drive 427, a magnetic disk drive 428, *e.g.*, to read from or write to a removable disk 429, and an optical disk drive 430, *e.g.*, for reading a CD-ROM disk 431 or to read from or write to other optical media. The hard disk drive 427, magnetic disk drive 428, and optical disk drive 430 are connected to the system bus 423 by a hard disk drive interface 432, a magnetic disk drive interface 433, and an optical drive interface 434, respectively. The drives and their associated computer-readable media provide nonvolatile storage of data, data structures, computer-executable instructions, etc. for the server computer 420. Although the description of computer-readable media above refers to a hard disk, a removable magnetic disk and a CD, it should be appreciated by those skilled

in the art that other types of media which are readable by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, and the like, may also be used in the exemplary operating environment, and further that any such media may contain computer-executable instructions for performing the methods of the present invention.

5           A number of program modules may be stored in the drives and RAM 425, including an operating system 435, one or more application programs 436, other program modules 437, and program data 438. The operating system 435 in the illustrated computer may be a Microsoft operating system (*e.g.*, Windows NT operating system). It is to be appreciated that other operating systems may be employed such as UNIX for example.

10           A user may enter commands and information into the server computer 420 through a keyboard 440 and a pointing device, such as a mouse 442. Other input devices (not shown) may include a microphone, a joystick, a game pad, a satellite dish, a scanner, or the like. These and other input devices are often connected to the processing unit 421 through a serial port interface 446 that is coupled to the system bus, but may be connected by other interfaces,  
15           such as a parallel port, a game port or a universal serial bus (USB). A monitor 447 or other type of display device is also connected to the system bus 423 via an interface, such as a video adapter 448. In addition to the monitor, computers typically include other peripheral output devices (not shown), such as speakers and printers.

20           The computer 420 may operate in a networked environment using logical connections to one or more remote computers, such as a remote client computer 449. The remote computer 449 may be a workstation, a server computer, a router, a peer device or other common network node, and typically includes many or all of the elements described relative to the server computer 420, although only a memory storage device 450 is illustrated in FIG. 8. The logical connections depicted in FIG. 8 include a local area network (LAN) 451 and a  
25           wide area network (WAN) 452. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

          When employed in a LAN networking environment, the server computer 420 may be connected to the local network 451 through a network interface or adapter 453. When utilized in a WAN networking environment, the server computer 420 generally may include a

modem 454, and/or is connected to a communications server on the LAN, and/or has other means for establishing communications over the wide area network 452, such as the Internet. The modem 454, which may be internal or external, may be connected to the system bus 423 *via* the serial port interface 446. In a networked environment, program modules depicted  
5 relative to the computer 420, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

In accordance with the practices of persons skilled in the art of computer programming, the present invention has been described with reference to acts and symbolic  
10 representations of operations that are performed by a computer, such as the computer 420, unless otherwise indicated. Such acts and operations are sometimes referred to as being computer-executed. It will be appreciated that the acts and symbolically represented operations include the manipulation by the processing unit 421 of electrical signals  
15 representing data bits which causes a resulting transformation or reduction of the electrical signal representation, and the maintenance of data bits at memory locations in the memory system (including the system memory 422, hard drive 427, floppy disks 429, and CD-ROM 431) to thereby reconfigure or otherwise alter the computer system's operation, as well as  
20 other processing of signals. The memory locations wherein such data bits are maintained are physical locations that have particular electrical, magnetic, or optical properties corresponding to the data bits.

What has been described above are preferred aspects of the present invention. It is, of course, not possible to describe every conceivable combination of components or methodologies for purposes of describing the present invention, but one of ordinary skill in  
25 the art will recognize that many further combinations and permutations of the present invention are possible. Accordingly, the present invention is intended to embrace all such alterations, modifications and variations that fall within the spirit and scope of the appended claims.



### **Claims**

What is claimed is:

1. A system for gathering and aggregating operational metrics of a plurality of members configured as an entity, comprising:

at least one of the plurality of members having a gathering and aggregation system adapted to gather operational metric data from the plurality of members and aggregate the operational metric data into a unified result set.

2. The system of claim 1, wherein each of the plurality of members have a gathering and aggregation system such that the unified result set can be gathered and aggregated from any of the plurality of members.

3. The system of claim 1, further comprising a metric monitor system for each of the plurality of members, the metric monitoring system monitoring member specific metrics and employing a time aggregation component to aggregate member specific metrics over time.

4. The system of claim 3, the time aggregation component being further operable to aggregate member specific performance metrics data into data of larger time periods and larger resolutions.

5. The system of claim 4, the time aggregation component aggregating member specific performance metrics data into data of larger time periods and larger resolutions by taking one of an average, a minimum, a maximum, a last and a weighted average of performance metrics data of a first time period and first resolution to evaluate performance metric data of a consecutive time period and consecutive resolution.

6. The system of claim 1, the gathering and aggregation system further comprising a performance metric aggregation component adapted to gather and aggregate performance metric data values of a particular time period and resolution from the plurality of members based on a time period requested by a requestor.

7. The system of claim 6, the performance metric aggregation component being operable to aggregate data performance values having similar data times to form a unified result set over the particular time period and time resolution.

8. The system of claim 7, the performance metric aggregation component aggregating member specific performance metrics data into a unified result set by evaluating a single data value for data points of similar data times by taking one of an average, a minimum, a maximum, a last and a weighted average data of similar data times.

9. The system of claim 1, the plurality of members logging member specific operational metrics to a data store corresponding to that particular member.

10. The system of claim 1, the gathering and aggregation system being further adapted to receive a request from a requestor for operational metric information for the entity and return the unified result set back to the requestor.

11. The system of claim 10, the requestor being one of an external process, an internal process, an external consumer, a user interface and another entity.

12. The system of claim 1, at least one of the plurality of members being configurable to receive a configuration setting defining the operational metric information to be logged and replicating the configuration setting to the plurality of members.

13. The system of claim 12, wherein any of the plurality of members are

configurable to receive a configuration setting defining the operational metric information to be logged.

14. The system of claim 1, the operational metric data being at least one of performance metric data, event metric data, status metric data and health metric data.

15. The system of claim 1, the gathering and aggregation system being operable to aggregate valid operational metric data and compensate for invalid operational metric data.

16. The system of claim 1, the gathering and aggregation system being further adapted to provide a unified result set of operational metric data for a single member.

17. A system for gathering and aggregating operational metrics of a plurality of members configured as an entity, comprising:

at least one of the plurality of members having a gathering and aggregation system adapted to gather operational event metric data from the plurality of members and coalesce the operational event metric data into a unified result set.

18. The system of claim 17, the gathering and aggregation system being further adapted to receive a request from a requestor for operational event metric data for the entity and return a portion of the unified result set back to the requestor.

19. The system of claim 18, each event having a unique identifier, the gathering and aggregation system employing at least one of the unique identifier of a first and a last event of the portion to provide one of a previous and subsequent portion of the unified result set back to the requestor upon request by the requestor and a specific unique identifier for additional information on a specific event upon request by the requestor.

20. The system of claim 19, the unique identifier comprising an event identifier,

an event source and an event time.

21. The system of claim 18, further comprising event filters for filtering out events not to be returned to the requestor.

22. A system for gathering and aggregating operational metrics of a plurality of members configured as an entity to determine a health state of the entity, comprising:

at least one of the plurality of members having a gathering and aggregation system adapted to gather operational metric data from the plurality of members and aggregate the operational metric data into at least one of a single result state based on a set of predefined rules.

23. The system of claim 22, each of the plurality of members determining a member health state based on a member specific health rule set, the gathering and aggregation system gathering the member health states from the plurality of members and applying a health entity aggregation rule set to determine a health state of the entity.

24. The system of claim 22, the gathering and aggregation system aggregating the operational metric data from the plurality of members into a unified result set and applying a health member aggregation rule set to determine a health state of the entity.

25. A method for gathering and aggregating operational metrics of a plurality of members configured as an entity, comprising the step of:

querying an operational metric from a plurality of members; and  
aggregating the operational metric from the plurality of members to form a unified result set.

26. The method of claim 25, further comprising the steps of receiving a request from a requestor for an operational metric for the entity prior to the step of querying and

returning the unified result step after the step of aggregating.

27. The method of claim 25, further comprising the step aggregating the operational metric over time at each of the plurality of members prior to the step of querying.

28. The method of claim 27, the step aggregating the operational metric over time at each of the plurality of members comprising the step of aggregating the operational metric into data of at least one larger time period and larger time resolution.

29. The method of claim 28, the operational metric being an operational performance metric and the step of aggregating the operational metric from the plurality of members to form a unified result set comprising aggregating data performance values having similar data times to form a unified result set.

30. The method of claim 25, further comprising a step of setting a configuration at one of the plurality of members defining the operational metric data to be logged at each of the plurality of members and replicating the configuration to each of the plurality of members.

31. The method of claim 25, the operational metric being an operational event metric and the step of aggregating the operational metric from the plurality of members to form a unified result set comprising the step of coalescing events from the plurality of members to form a unified result set.

32. The method of claim 31, further comprising the steps of receiving a request from a requestor for events for the entity prior to the step of querying and returning a portion of the unified result step after the step of aggregating.

33. The method of claim 32, further comprising the step of assigning each event a

unique identifier and the step of requesting and returning additional event data utilizing the unique identifier of at least one event.

34. The method of claim 32, further comprising the step of filtering out events not to be returned to the requestor.

35. The method of claim 25, the operational metric being an operational health metric and the step of aggregating the operational metric from the plurality of members to form a unified result set comprising the step of aggregating the operational data into at least one of a single result state based on a set of predefined rules.

36. The method of claim 35, further comprising the steps of determining a member health state for each member based on a member specific health rule set, querying each of the member health states from the plurality of members and applying a health entity aggregation rule set to determine a health state of the entity.

37. The method of claim 35, further comprising the step of applying a health member aggregation rule set to the unified result set to determine a health state of the entity.

38. A computer-readable medium having computer-executable instructions for performing the steps of claim 25.

39. A system for gathering and aggregating operational metrics of a plurality of members configured as an entity, comprising:

means for querying an operational metric from a plurality of members;

means for aggregating the operational metric from the plurality of members to form a unified result set.

40. The system of claim 39, further comprising means for receiving a request from

a requestor for an operational metric for the entity and means for returning the unified result set.

41. The system of claim 39, further comprising means for setting a configuration at one of the plurality of members defining the operational metric data to be logged at each of the plurality of members and means for replicating the configuration to each of the plurality of members.

42. The system of claim 39, the operational metric being at least one of a performance metric, an event metric, a status metric and a health metric.

DECLASSIFIED


### **Abstract of the Invention**

A system and method is provided for gathering and aggregating operational metrics (e.g., performance metrics, process events, health monitor state, server state) for a plurality of members as a single entity and for a plurality of entities as a singular entity. The system and method provides for operational metrics of members and entities to be aggregated and retrieved as a single result set, such that entity wide operational metrics can be acquired, monitored and displayed as a single entity.


10

15





**Fig. 1**



**Fig. 2**






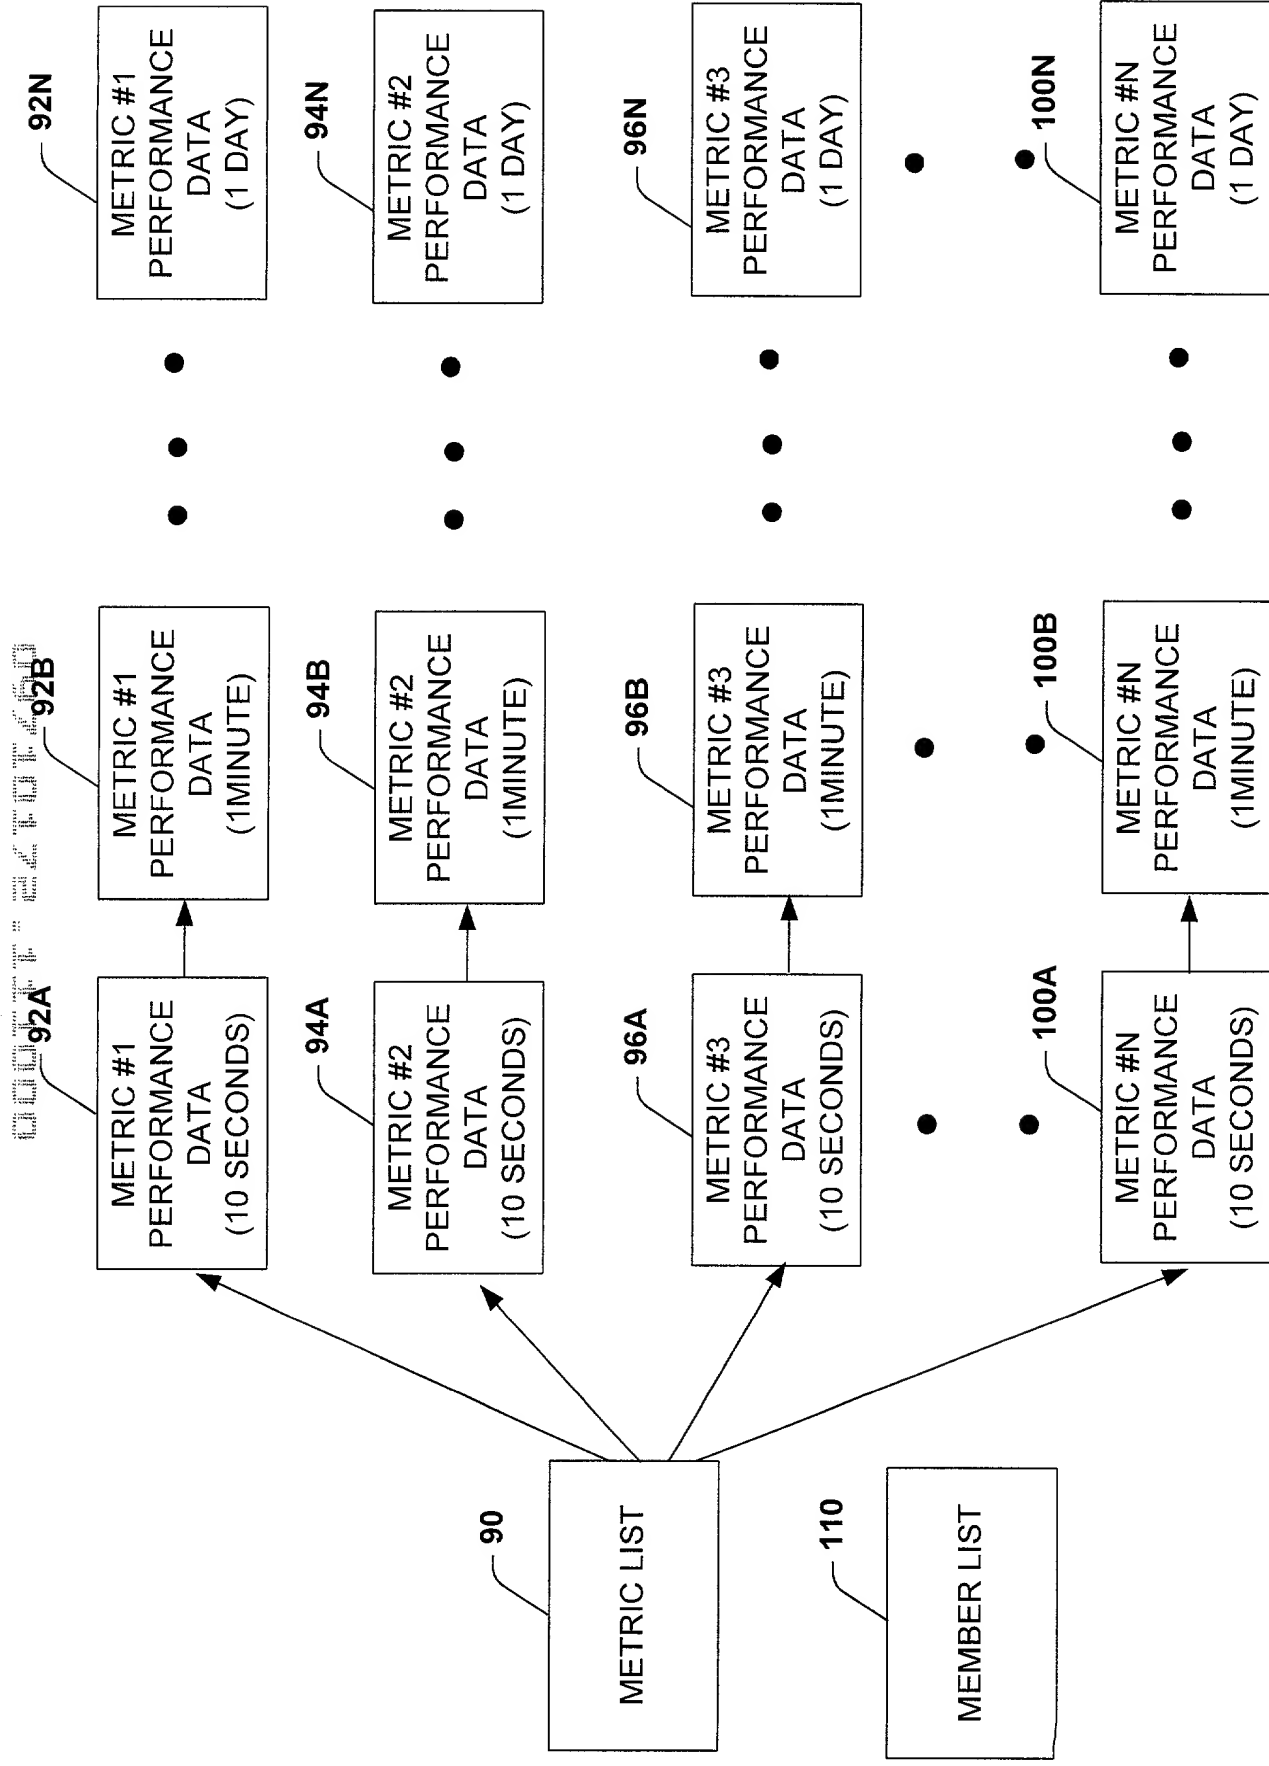
Fig. 3




**Fig. 4a**



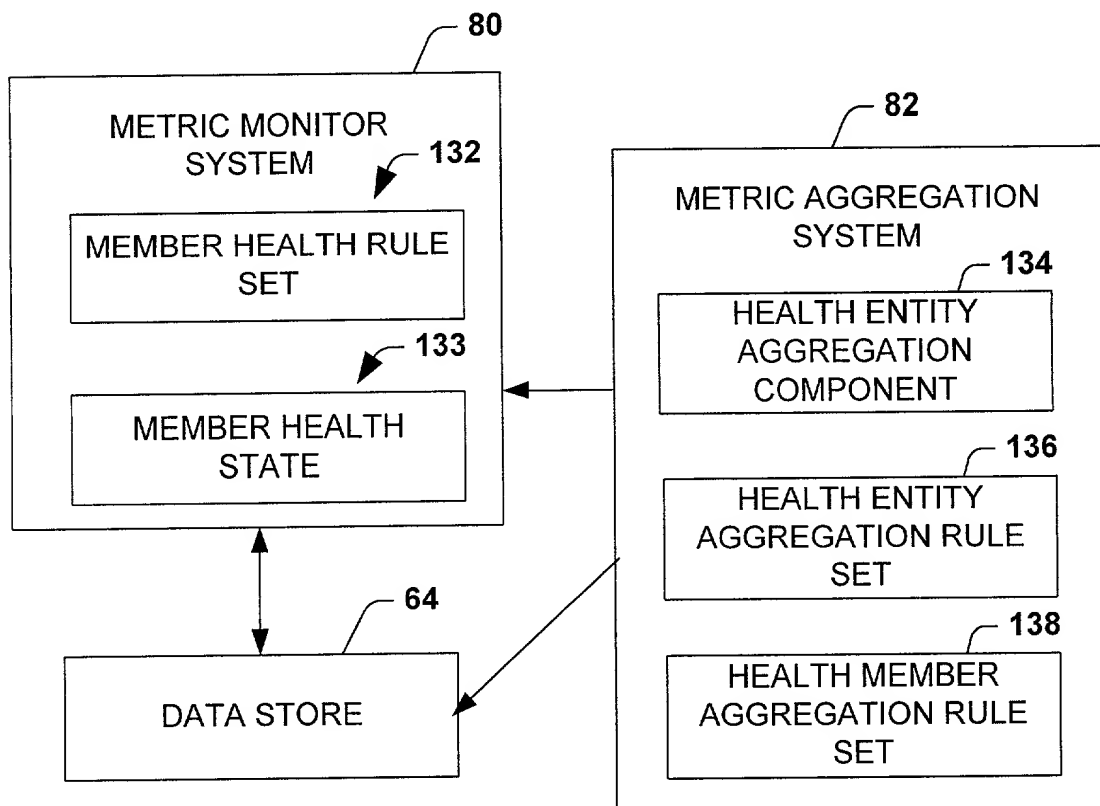
**Fig. 4b**



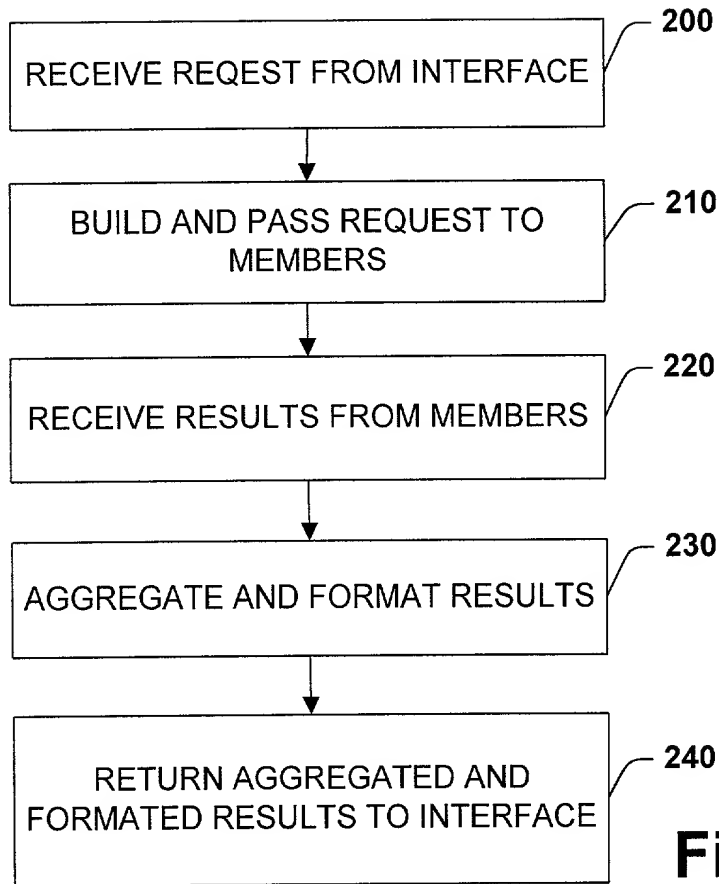
**Fig. 5**



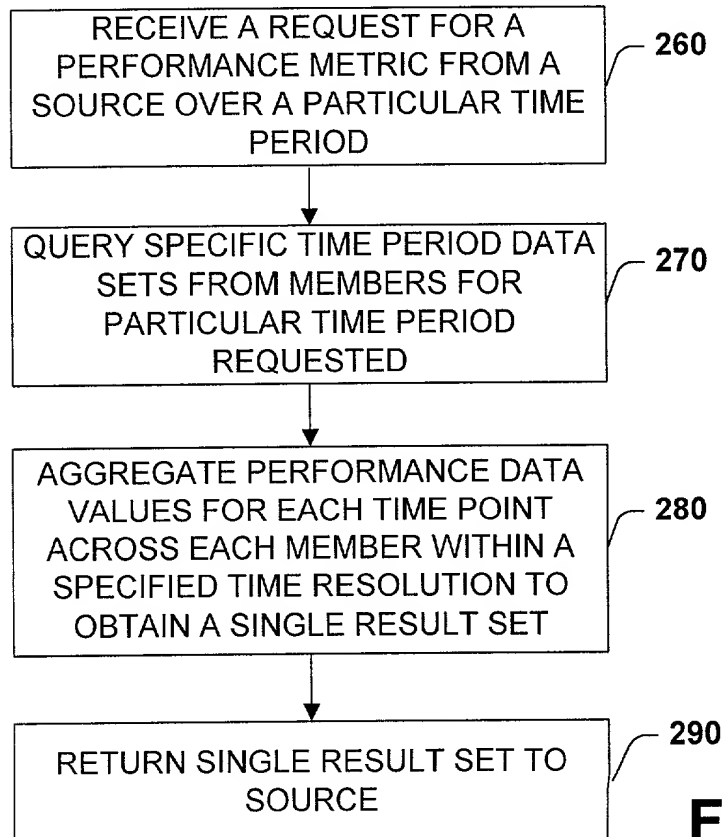
**Fig. 6a**



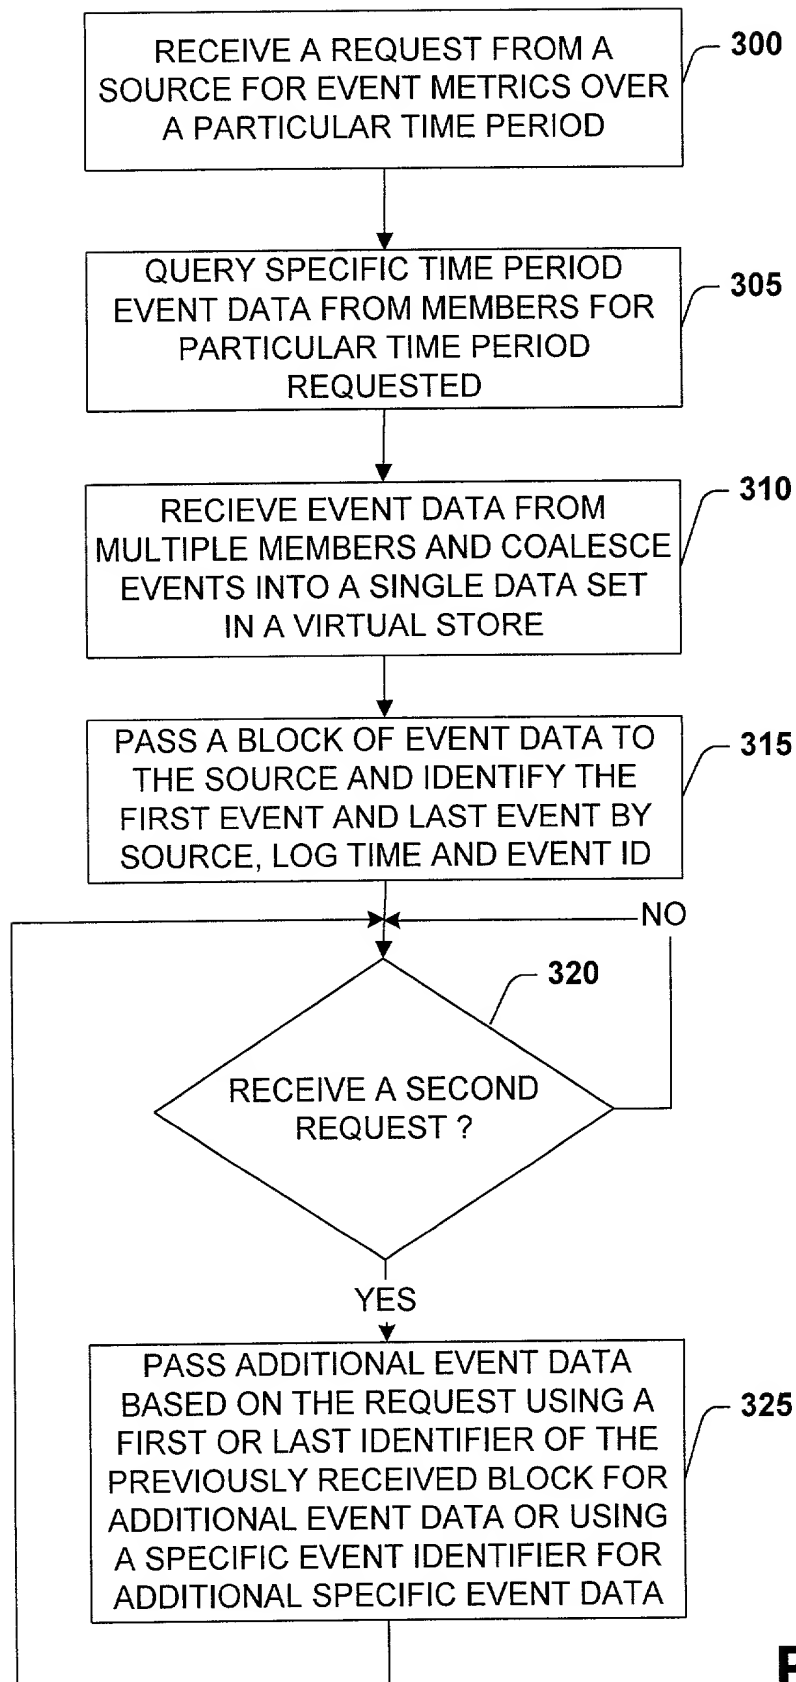
**Fig. 6b**



**Fig. 7a**

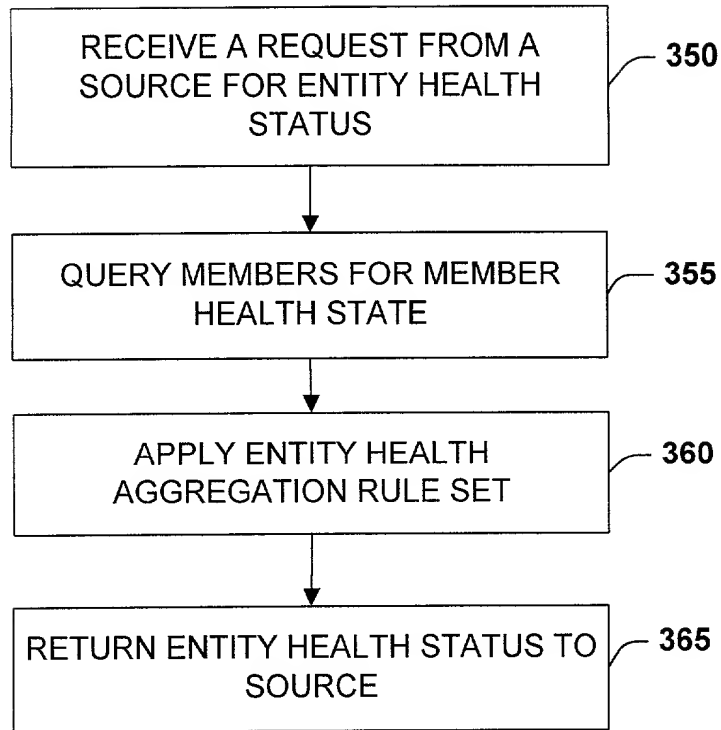


**Fig. 7b**

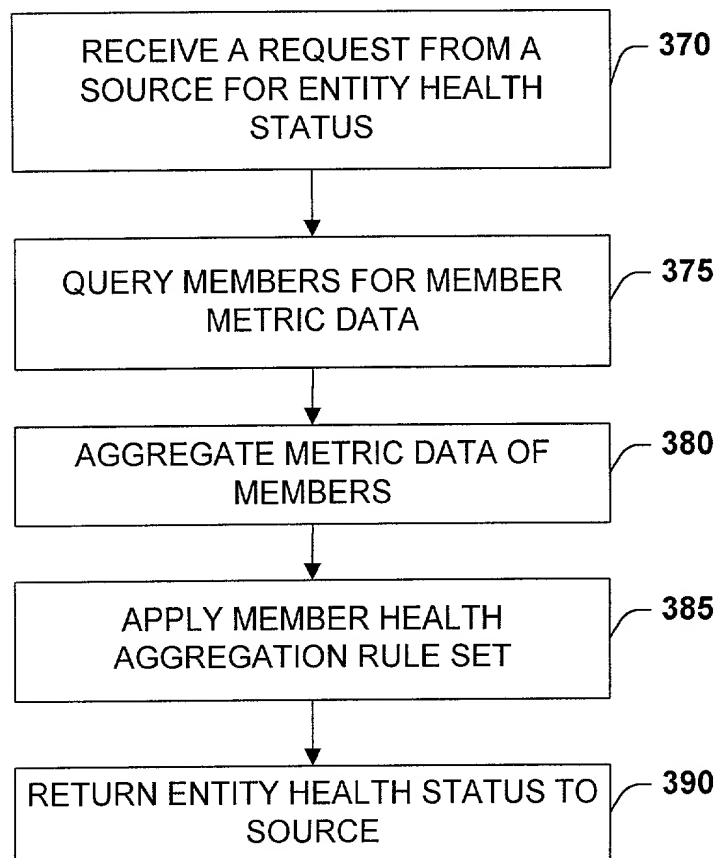


**Fig. 7c**

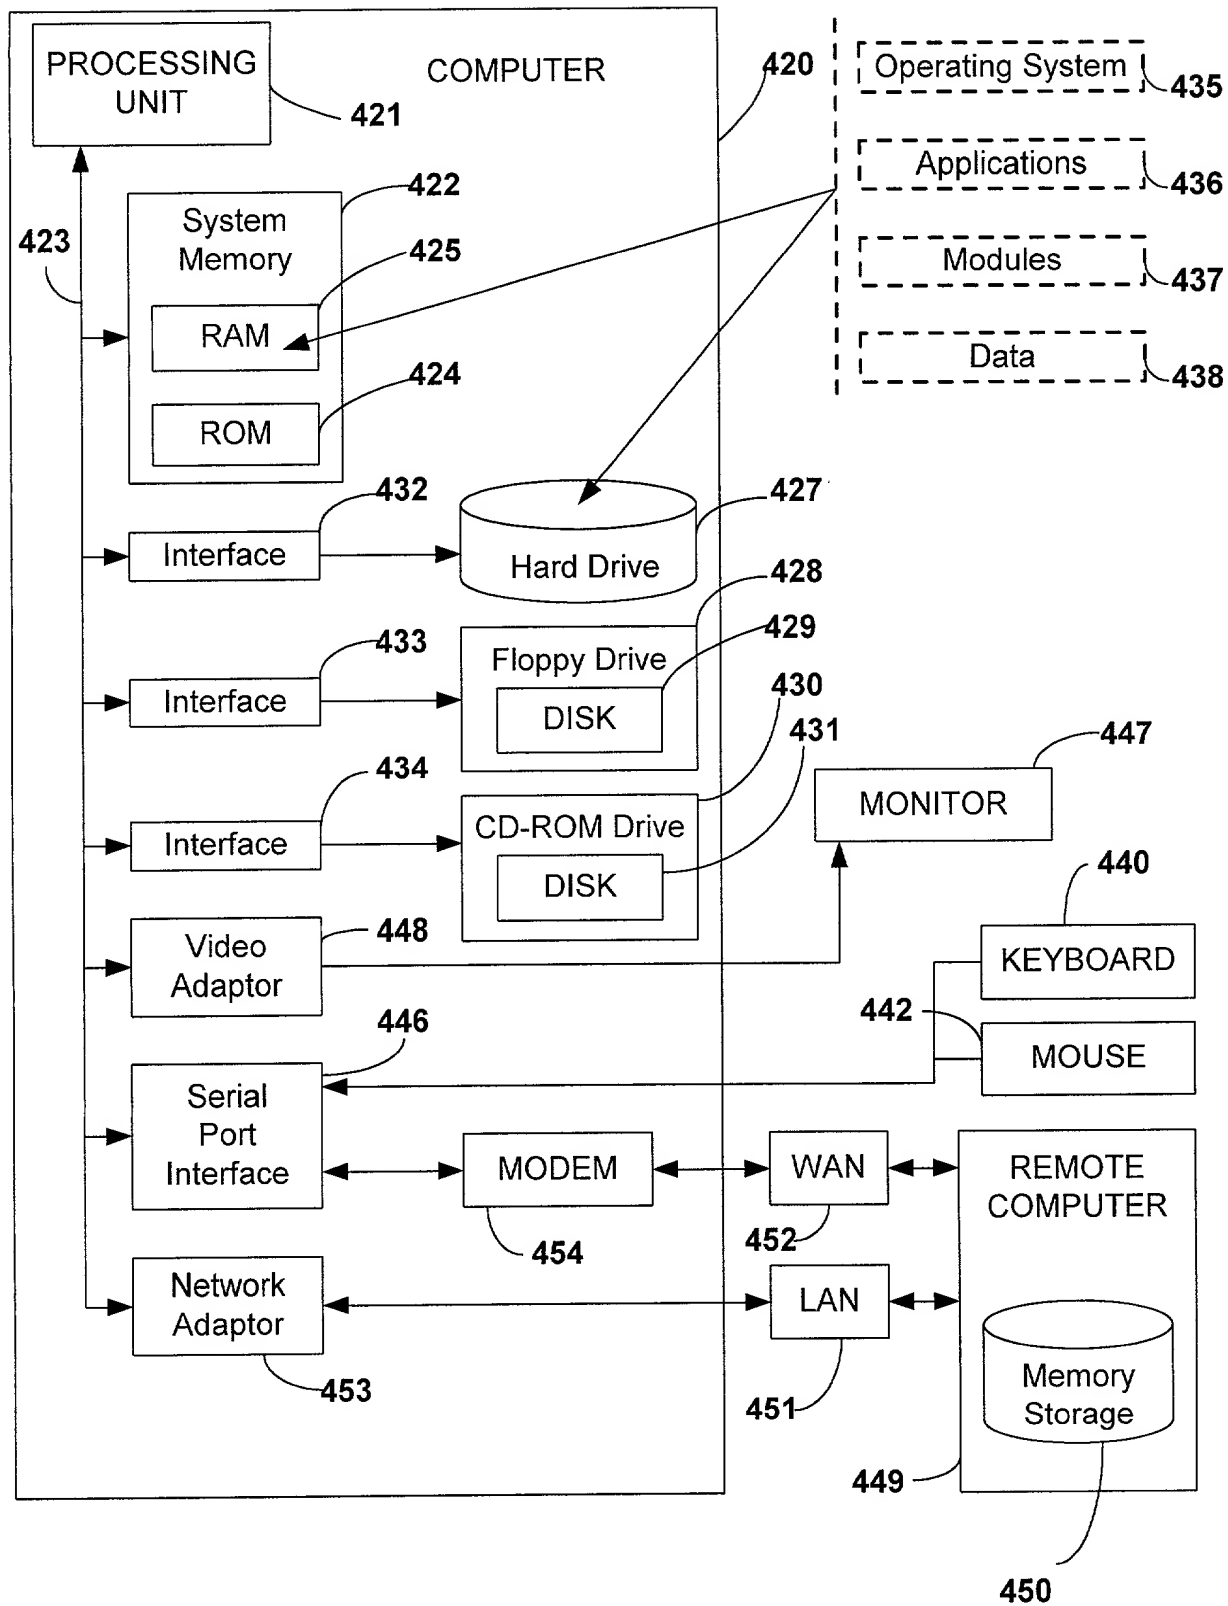




**Fig. 7d**



**Fig. 7e**



**Fig. 8**

**COMBINED DECLARATION AND POWER OF ATTORNEY**  
(ORIGINAL, DESIGN, NATIONAL STAGE OF PCT)

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name, I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention

entitled: **DISTRIBUTED DATA GATHERING AND AGGREGATION AGENT**

the specification of which:

- (a) ☒ is attached hereto.  
 (b) \_\_\_ was filed on \_\_\_\_\_ as Serial No. \_\_\_\_\_ or  
 Express Mail No. \_\_\_\_\_, as Serial No. not yet known, and was amended on  
 (if applicable).  
 (c) \_\_\_ was described and claimed in PCT International Application No. \_\_\_\_\_ filed  
 on \_\_\_\_\_ and amended under PCT Article 19 on \_\_\_\_\_ (if any).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to patentability in accordance with Title 37, Code of Federal Regulations '1.56(a).

**PRIORITY CLAIM**

I hereby claim foreign priority benefits under Title 35, United States Code, '119 of any foreign application(s) for patent or inventor's certificate or of any PCT international application(s) designating at least one country other than the United States of America listed below and have also identified below any foreign application(s) for patent or inventor's certificate or any PCT international application(s) designating at least one country other than the United States of America filed by me on the same subject matter having a filing date before that of the application(s) of which priority is claimed.

- (d) ☒ no such applications have been filed.  
 (e) \_\_\_ such applications have been filed as follows.

**EARLIEST FOREIGN APPLICATION(S), IF ANY FILED WITHIN 12 MONTHS  
(6 MONTHS FOR DESIGN) PRIOR TO THIS U.S. APPLICATION**

COUNTRY	APPLICATION NUMBER	DATE OF FILING (day, month, year)	PRIORITY CLAIMED UNDER 35, USC 119
_____	_____	_____	___ Yes ___ No
_____	_____	_____	___ Yes ___ No
_____	_____	_____	___ Yes ___ No

**ALL FOREIGN APPLICATION(S), IF ANY FILED MORE THAN 12 MONTHS  
(6 MONTHS FOR DESIGN) PRIOR TO THIS U.S. APPLICATION**

\_\_\_\_\_  
 \_\_\_\_\_

**POWER OF ATTORNEY**

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (List name and registration number)

Himanshu S. Amin, Reg. No. 40,894; Gregory Turocy, Reg. No. 36, 952;  
Christopher P. Harris, Reg. No. 43,660; Eric M. Highman,  
Reg. No. 43,672; and Gary J. Pitzer, Reg. No. 39,334.

Katie E. Sako, Reg. No. 32,628 and Daniel D. Crouse, Reg. No. 32,022.

The undersigned to this declaration and power of attorney hereby authorizes the U.S. attorney(s) named herein to accept and follow instructions from:

Name(s) of authorized representative(s) \_\_\_\_\_  
Address \_\_\_\_\_

as to any actions to be taken in the Patent and Trademark Office regarding this application without direct communication between the U.S. attorney(s) and the undersigned. In the event of a change in the person(s) from whom instructions may be taken, the U.S. attorney(s) will be so notified by the undersigned.

**Send Correspondence To:**

Himanshu S. Amin  
AMIN, ESCHWEILER & TUROCY, LLP  
24<sup>TH</sup> Floor, National City Center  
1900 East 9<sup>TH</sup> Street  
Cleveland, Ohio 44114

**Direct Telephone Calls To:**  
(name and telephone number)

Himanshu S. Amin  
(216) 696-8730

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with knowledge that willful false statements and the like are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued therein.

Full name of sole or first inventor, if any: Federico Garcea  
Inventor's signature: *Federico Garcea*  
Date: 11/6/2000 Country of Citizenship: Italy  
Residence: Seattle, Washington  
Post Office Address: 119 NW 75<sup>th</sup> Street  
Seattle, Washington 98117

Full name of second joint inventor, if any: Michael S. Murstein  
Inventor's signature: *Michael S. Murstein*  
Date: 11/6/2000 Country of Citizenship: U.S.  
Residence: Kirkland, Washington  
Post Office Address: 6511 126<sup>th</sup> Avenue NE  
Kirkland, Washington 98033

CHECK FOR ANY OF THE FOLLOWING ADDED PAGE(S) WHICH  
FORM A PART OF THIS DECLARATION

X Signature for third and subsequent joint inventors. Number of pages added 1.

Full name of **third or joint inventor**, if any: Roger W. Sprague  
Inventor's signature: *Roger W. Sprague*  
Date: 11/6/2000 Country of Citizenship: U.S.  
Residence: Redmond Washington  
Post Office Address: 7219 151<sup>st</sup> Street  
Redmond, Washington 98052

Full name of **fourth joint inventor**, if any: Alexander M. Sutton  
Inventor's signature: *Alexander M. Sutton*  
Date: 11/6/2000 Country of Citizenship: U.S.  
Residence: Seattle, Washington  
Post Office Address: 920 N 34<sup>th</sup> Street, Unit 50  
Seattle, Washington 98103

Full name of **fifth joint inventor**, if any: Michael W. Thomas  
Inventor's signature: *Michael W. Thomas*  
Date: 11/6/2000 Country of Citizenship: U.S.  
Residence: Bellevue, Washington  
Post Office Address: 1134 166<sup>th</sup> Avenue SE  
Bellevue, Washington 98008

Full name of **sixth joint inventor**, if any: Giedrius Zizys  
Inventor's signature: *Giedrius Zizys*  
Date: 11/6/2000 Country of Citizenship: Lithuania  
Residence: Redmond, Washington  
Post Office Address: 13680 NE 69<sup>th</sup> Street, Apt. 637  
Redmond, Washington 98052

X This declaration ends with this page.